

AD-A160 233

FOUNDATIONS OF KNOWLEDGE FOR DISTRIBUTED SYSTEMS(U)
YALE UNIV NEW HAVEN CT DEPT OF COMPUTER SCIENCE
M J FISHER ET AL. SEP 85 TR-426 N00014-82-K-0154

1/1

UNCLASSIFIED


F/G 9/2

NL

END

FORMED


DNF



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

AD-A160 233

12



FOUNDATIONS OF KNOWLEDGE FOR DISTRIBUTED SYSTEMS

Michael J. Fischer and Neil Immerman

YALEU/DCS/TR 426
September, 1985

DTIC FILE COPY

This document has been approved
for public release and sale; its
distribution is unlimited.

YALE UNIVERSITY
DEPARTMENT OF COMPUTER SCIENCE

DTIC
ELECTE
S OCT 16 1985 D
E

85 10 15 075

FOUNDATIONS OF KNOWLEDGE FOR DISTRIBUTED SYSTEMS

Michael J. Fischer and Neil Immerman

YALEU/DCS/TR 426
September, 1985

This document has been approved
for public release and ordering
distribution is unlimited.

DTIC
UNCLASSIFIED
OCT 13 1985
E

| REPORT DOCUMENTATION PAGE | | READ INSTRUCTIONS BEFORE COMPLETING FORM |
|--|--------------------------------------|---|
| 1. REPORT NUMBER 426 | 2. GOVT ACCESSION NO. AD-A160 233 | 3. RECIPIENT'S CATALOG NUMBER |
| 4. TITLE (and Subtitle) FOUNDATIONS OF KNOWLEDGE FOR DISTRIBUTED SYSTEMS | | 5. TYPE OF REPORT & PERIOD COVERED Technical Report |
| | | 6. PERFORMING ORG. REPORT NUMBER |
| 7. AUTHOR(s) Michael J. Fischer and Neil Immerman | | 8. CONTRACT OR GRANT NUMBER(s) NSF: DCR-8405478 and ONR: N00014-82-K-0154 |
| 9. PERFORMING ORGANIZATION NAME AND ADDRESS Department of Computer Science/ Yale University Dunham Lab/ 10 Hillhouse Avenue New Haven, Connecticut 06520 | | 10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS |
| 11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research 800 N. Quincy Arlington, Virginia 22217 | | 12. REPORT DATE September, 1985 |
| | | 13. NUMBER OF PAGES 11 |
| 14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) | | 15. SECURITY CLASS. (of this report) Unclassified |
| | | 15a. DECLASSIFICATION/DOWNGRADING SCHEDULE |
| 16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distributed unlimited. | | |
| 17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) | | |
| 18. SUPPLEMENTARY NOTES | | |
| 19. KEY WORDS (Continue on reverse side if necessary and identify by block number) knowledge protocol common knowledge logic distributed system formal model foundations | | |
| 20. ABSTRACT (Continue on reverse side if necessary and identify by block number) General and precise definitions are given of distributed protocol and of various notions of knowledge and common knowledge in distributed systems. Whether or not a process in a distributed protocol can obtain new common knowledge during execution is shown to depend on which notion of common knowledge chosen, showing that the problem of capturing intuitive concepts about knowledge is more subtle than was previously believed. | | |

Foundations of Knowledge for Distributed Systems

Michael J. Fischer* and Neil Immerman†

Computer Science Department
Yale University
New Haven, CT 06520

September 19, 1985

1 Introduction

In [HM84], Halpern and Moses present an interesting discussion of knowledge and common knowledge for distributed systems. They argue that while common knowledge is desirable, it is unattainable in certain settings. They suggest a hierarchy of weakened versions of common knowledge and discuss when these can be achieved.

One difficulty with the Halpern and Moses paper is that it is informal and the concepts dealt with are not rigorously defined. Given one interpretation, their theorems are true as claimed, but given another, the opposite occurs, as we show with appropriate counterexamples. Thus, their theorems are not wrong in spirit, but the concepts are subtle and the terms must be more carefully defined.

In the main body of this paper, we give quite general and precise definitions of distributed protocol, knowledge and common knowledge. We also provide precise settings and rigorous proofs for many of the results in [HM84]. We then propose some alternate definitions in which the same results become false.

The author Our desire is to develop a way to design clear distributed algorithms and write clear proofs about them. We believe we have provided a solid base for future work in this area.

*Research supported in part by the National Science Foundation under Grant Number DCR-8405478 and by the Office of Naval Research under Contract Number N00014-82-K-0154.

†Research supported by an NSF postdoctoral fellowship.

*Additional keywords:
distributed data processing*



| | |
|--------------------|-------------------------------------|
| Accession For | |
| NTIS GRA&I | <input checked="" type="checkbox"/> |
| DTIC TAB | <input type="checkbox"/> |
| Unannounced | <input type="checkbox"/> |
| Justification | |
| By | |
| Distribution/ | |
| Availability Codes | |
| Avail and/or | |
| Dist | Special |
| A-1 | |

2 Definitions

Definition 2.1 A distributed protocol,

$$P = (n, Q, I, \tau),$$

consists of a number n of participants, a set Q of local states, a set $I \subseteq Q^n$ of initial global states, and a next move relation $\tau \subseteq Q^n \times Q^n$ on global states.

Our definition of protocol is certainly simple and precise. Let us first argue that it is sufficiently general. Anything we would be willing to call a distributed system can be broken up into a finite number of logical entities. Each such entity must be in some total configuration that we are calling its state. Furthermore the states of all the components combined should determine the entire state of the system and thus which global states can be next entered.

It is easy to see for example that our model of distributed system is a generalization of the shared variable model of Lynch and Fischer [LF81]. In that model the components consist of shared variables and processors. Each action involves exactly one processor and one shared variable.

Similarly our model includes synchronous protocols in which every processor sends a message to every other during each round. One way to model this is to specify that the set of possible states is of the form $Q = M^n$, i.e. each processor's total configuration consists of an n -tuple. We can specify that the i^{th} entry of j 's state is the value of the message sent from i to j during the previous round. This can be done as follows: for all processors i, j , and for all global states p, q, r, s , if (p, q) and (r, s) are in τ and if processor i has the same state in p as in r , then the i^{th} component of processor j 's state is the same in q as in s .

The sense in which our model could be too general is that we allow any transition relation τ . Of course for certain applications we can make appropriate restrictions. We have already seen that we can restrict our attention to processors which communicate with shared variables, or to synchronous message passing protocols. Similarly, instead of letting each processor's transitions be perfectly general, we can restrict our attention to processors with specified computing power, e.g. finite automaton, polynomial time Turing machine, etc.

For any protocol P , let R_P be the τ -reachable global states of P , that is, the set of all global states we can reach by starting in I and taking any number of τ steps. For $p \in Q^n$ a global state and $1 \leq i \leq n$, we write $(p)_i$ to denote the i^{th} component of p .

For any two states $p, q \in R_P$ and any participant i , we will use the notation $p \stackrel{i}{\sim} q$ to mean that $(p)_i = (q)_i$, i.e. they are indistinguishable from i 's point of view. Obviously each $\stackrel{i}{\sim}$ is an equivalence relation. For any reachable global state p , define the i -neighborhood of p as follows:

$$N_i(p) = \{q \mid q \stackrel{i}{\sim} p\}$$

If we are in state p , then all i "knows" is that we are in $N_i(p)$. Therefore, for any sentence α ,¹ it is natural to make the following definition of $K_i\alpha$, which we read as " i knows α ":

$$\langle P, p \rangle \models K_i\alpha \quad \equiv \quad \forall q \in N_i(p) (\langle P, q \rangle \models \alpha)$$

Intuitively i knows α just if α is true in all the worlds which are indistinguishable by i from the current world.

It is convenient to picture a protocol P as a graph with nodes consisting of all the elements of R_P . There is a directed edge labelled τ from p to q just if $\langle p, q \rangle \in \tau$. Furthermore there is an undirected edge labelled ' i ' between p and q just if $p \sim_i q$.

Let $G \subseteq \{1, \dots, n\}$ be a group of participants in a protocol. For any $p \in R_P$, define the G -neighborhood of p as follows:

$$N_G(p) = \{q \mid (\exists r \geq 0)(\exists i_1, \dots, i_r \in G)(\exists p_1, \dots, p_{r-1})[p \stackrel{i_1}{\sim} p_1 \stackrel{i_2}{\sim} p_2 \dots p_{r-1} \stackrel{i_r}{\sim} q]\}$$

This generalizes our previous definition since $N_i(p) = N_{\{i\}}(p)$.

Analogously to our definition of $K_i\alpha$, we define $C_G\alpha$, which we read, "it is common knowledge among the members of G that α ":

$$\langle P, p \rangle \models C_G\alpha \quad \equiv \quad \forall q \in N_G(p) (\langle P, q \rangle \models \alpha)$$

We write C for C_G and N for N_G in the special case that G includes all participants.

The next result shows that $C\alpha$ coincides with the intuitive definition current in the literature. (See for example [HM84].)

Theorem 2.2 *The following two statements are equivalent:*

1. $\langle P, p \rangle \models C_G\alpha$.
2. $(\forall r \geq 0)(\forall i_1, \dots, i_r \in G)(\langle P, p \rangle \models K_{i_1}K_{i_2} \dots K_{i_r}\alpha)$.

Proof

(1 \Rightarrow 2): For any β , we have $C_G\beta \rightarrow \beta$ since $p \in N_G(p)$. Thus, it suffices to show that for any β , if $\langle P, p \rangle \models C_G\beta$, then for all $i \in G$, $\langle P, p \rangle \models C_GK_i\beta$. This is clear because if $q \in N_G(p)$ and $q' \in N_i(q)$ then $q' \in N_G(p)$; hence $\langle P, q \rangle \models K_i\beta$. Since $K_i\beta$ holds for all $q \in N_G(p)$, it is common knowledge in G at p , as desired.

(2 \Rightarrow 1): Suppose that $\langle P, p \rangle \not\models C_G\alpha$. It follows that there is a $q \in N_G(p)$ such that $\langle P, q \rangle \models \neg\alpha$. Let $i_1, \dots, i_r \in G$ be such that there exists p_1, \dots, p_{r-1} with $p \stackrel{i_1}{\sim} p_1 \stackrel{i_2}{\sim} p_2 \dots p_{r-1} \stackrel{i_r}{\sim} q$. It follows that $\langle P, p \rangle \models \neg K_{i_1}K_{i_2} \dots K_{i_r}\alpha$. ■

¹We have intentionally left the logical language unspecified from which the sentence α is drawn, for all that we require is that it be possible to interpret α at the pair $\langle P, p \rangle$.

We conclude this section with two nontrivial examples of protocols, one asynchronous and the other synchronous. These protocols will be frequently referred to in the remainder of the paper.

Let $A = \langle n^2, Q_A, I_A, \tau_A \rangle$ be an asynchronous message passing protocol defined as follows: The first n participants of A are the processors a_1, \dots, a_n ; the remaining $(n-1)n$ participants are buffers. For a global state p , we abuse our previous notation slightly and write $(p)_{a_i}$ to denote the component corresponding to a_i and $(p)_{b_{i,j}}$ to denote the component corresponding to buffer $b_{i,j}$.

The set of possible local states of a buffer $b_{i,j}$, $i \neq j$, is $M \cup \{\lambda\}$, where M is a set of possible messages and λ is a special symbol denoting the null message. $b_{i,j} = m \in M$ indicates that the single message m was sent by i but not yet delivered to j . $b_{i,j} = \lambda$ indicates that no message is waiting.

The set of possible local states of a processor a_i is $(D \times M^{n-1} \times N)$. State $\langle d, m_1, \dots, m_{i-1}, m_{i+1}, \dots, m_n, r \rangle$ indicates that the processor is in internal state d at round r with pending messages $m_1, \dots, m_{i-1}, m_{i+1}, \dots, m_n$. If r is even, then the processor is in a 'send' state, waiting to place each $m_j \neq \lambda$ into buffer $b_{i,j}$. If r is odd, then the processor is in a 'receive' state waiting to fetch a message from $b_{j,i}$ for each j such that $m_j = \lambda$.

Thus, the complete set of local states Q_A is $(M \cup \{\lambda\}) \cup (D \times M^{n-1} \times N)$.

The transitions making up τ_A are of four kinds:

1. $\langle p, q \rangle \in \text{Send}_{i,j}$ if
 - $p \stackrel{\varepsilon}{\sim} q$ for all $c \notin \{a_i, b_{i,j}\}$;
 - $(p)_{b_{i,j}} = \lambda$;
 - $(q)_{b_{i,j}} = m_j \neq \lambda$;
 - $(p)_{a_i} = \langle d, \dots, m_{j-1}, m_j, \dots, 2k \rangle$;
 - $(q)_{a_i} = \langle d, \dots, m_{j-1}, \lambda, \dots, 2k \rangle$.
2. $\langle p, q \rangle \in \text{Receive}_{i,j}$ if
 - $p \stackrel{\varepsilon}{\sim} q$ for all $c \notin \{a_i, b_{j,i}\}$;
 - $(p)_{b_{j,i}} = m_j \neq \lambda$;
 - $(q)_{b_{j,i}} = \lambda$;
 - $(p)_{a_i} = \langle d, \dots, m_{j-1}, \lambda, \dots, 2k+1 \rangle$;
 - $(q)_{a_i} = \langle d, \dots, m_{j-1}, m_j, \dots, 2k+1 \rangle$.
3. $\langle p, q \rangle \in \text{Stop}_i$ if
 - $p \stackrel{\varepsilon}{\sim} q$ for all $c \neq a_i$;
 - $(p)_{a_i} = \langle d, \lambda, \dots, \lambda, 2k \rangle$;
 - $(q)_{a_i} = \langle d, \lambda, \dots, \lambda, 2k+1 \rangle$.

4. $\langle p, q \rangle \in \text{Start}_i$ if

- $p \sim q$ for all $c \neq a_i$;
- $\langle p \rangle_{a_i} = \langle d, m_1, \dots, m_n, 2k+1 \rangle$, where $m_j \neq \lambda$ for all $j \neq i$;
- $\langle q \rangle_{a_i} = \langle d', m'_1, \dots, m'_n, 2k+2 \rangle$, where $m'_j \neq \lambda$ for all $j \neq i$, and d' and m'_j are functions of $\langle d, \bar{m}, 2k+1 \rangle$.

Now we let τ_A consist of all the above transitions:

$$\tau_A = \bigcup_{i,j} \text{Send}_{i,j} \cup \text{Receive}_{i,j} \cup \text{Stop}_i \cup \text{Start}_i$$

Finally let I_A be some nonempty set of global states in which the state of every processor a_i has the form $\langle d, m_1, \dots, m_n, 0 \rangle$ with $m_j \neq \lambda$ for all $j \neq i$, and all the buffers are empty.

Our second example of a protocol is a synchronous version of A . Let $B = \langle n^2, Q_B, \tau_B, I_B \rangle$, where

$$Q_B = \{\lambda\} \cup (D \times M^{n-1} \times \{2r \mid r \in \mathbb{N}\})$$

Thus in B all buffers are empty and the local states of the a_i 's are the corresponding states from A at the beginning of a send phase. Let the transitions τ_B consist of all pairs $\langle p, q \rangle$ such that there exists a τ_A path in A from p to q such that none of the intermediate steps go through global states of B . Finally let $I_B = I_A$.

It is not hard to see that B is a synchronous version of A such that in each round all processes send $n-1$ messages and then receive $n-1$ messages.

3 Common Knowledge in Asynchronous Systems

Definition 3.1 We will call a protocol, $P = \langle n, Q, I, \tau \rangle$, totally asynchronous if for all $\langle p, q \rangle \in \tau$, p and q differ on at most two components.

The following theorem shows that in a totally asynchronous protocol, no new common knowledge can be achieved.

Theorem 3.2 Let P be a totally asynchronous protocol and G a set of at least three participants. Let p be any global state of P and let p_0 be an initial state from which p is reachable by a sequence of τ steps. Let α be any sentence in a logic for P . If $\langle P, p \rangle \models C_G \alpha$ then $\langle P, p_0 \rangle \models C_G \alpha$.

Proof We first show that if p is reachable from p_0 by a sequence of τ steps then $N_G(p) = N_G(p_0)$. It suffices to consider the case where $\langle p_0, p \rangle \in \tau$. By

the definition of a totally asynchronous protocol, p and p_0 must agree on all but at most two components. Thus, there exists a participant $j \in G$ with the same local state in p_0 as in p , i.e. $p_0 \stackrel{j}{\sim} p$. It follows that $p \in N_G(p_0)$ and thus $N_G(p) = N_G(p_0)$. The theorem now follows from the definition of common knowledge in G . ■

As an example, consider the protocol A discussed at the end of the last section. It is easy to check that A satisfies the definition of totally asynchronous and thus no new common knowledge can arise in A . By way of contrast, if we look at A 's cousin B , then one observes that all reachable global states in B have all processors in the same round. Thus, if two reachable global states are \sim equivalent for some i , $1 \leq i \leq n$, then they are both in the same round. It follows that if we let $G = \{a_1, \dots, a_n\}$ be the set of processors—i.e. we don't care what the buffers know—then at any round r , ' C_G (we're at round r)' holds, i.e. it is common knowledge in G that all processors are at round r .²

It would seem at first glance that the difficulty in achieving common knowledge has to do with the problem of reaching an arbitrary depth of K 's with only finitely many messages. We conclude this section with a look at finite state protocols where common knowledge is equivalent to a finite stack of K 's.

Theorem 3.3 Let $P = \langle n, Q, I, r \rangle$ be a finite state protocol, i.e. $|Q| < \infty$. For each i , let

$$Q_i = \{(q)_i \mid q \in R_P\}$$

Thus each processor is a $|Q_i|$ state automaton. Let $r = \min\{|Q_i| \mid 1 \leq i \leq n\}$. Let p be any global state and let α be any formula. Then the following are equivalent:

1. $\langle P, p \rangle \models C\alpha$.
2. For all $i_1, i_2, \dots, i_{2r-1}$, $(\langle P, p \rangle \models K_{i_1} \dots K_{i_{2r-1}} \alpha)$.

Proof

(1 \Rightarrow 2): By definition of C .

(2 \Rightarrow 1): Suppose that $\langle P, p \rangle \not\models C\alpha$. Then there must exist $q \in N(p)$ such that $\langle P, q \rangle \models \neg\alpha$. Consider a minimum length \sim chain from p to q :

$$p = p_0 \stackrel{i_1}{\sim} p_1 \stackrel{i_2}{\sim} p_2 \dots p_{s-1} \stackrel{i_s}{\sim} p_s = q$$

Note that no nonconsecutive pair p_j, p_k can agree on some component because if they did the chain could be shortened. It follows that in any given component each state appears at most twice. Therefore $s \leq 2r - 1$. It follows that

$$\langle P, p \rangle \models \neg K_{i_1} K_{i_2} \dots K_{i_{2r-1}} \alpha$$

²This assumes of course that our logical language is powerful enough to express the property 'we're at round r '.

Example 3.4 Consider the protocol $P_r = \langle 2, \{1, \dots, r+1\}, \{(1, 1)\}, \tau_r \rangle$ where,

$$\tau_r = \{((i, i), (i+1, i)) \mid 1 \leq i \leq r\} \cup \{((i+1, i), (i+1, i+1)) \mid 1 \leq i < r\}.$$

This protocol has the unique computation chain:

$$(1, 1), (2, 1), (2, 2), (3, 2), \dots, (r, r-1), (r, r), (r+1, r).$$

Furthermore, for all global states p and q we have $N(p) = N(q)$. Thus, for any α ,

$$\langle P_r, p \rangle \models C\alpha \Leftrightarrow \text{for all } q \in R_{P_r}, \langle P_r, q \rangle \models \alpha.$$

Let α say that processor 1 is not in state 1. Then $\langle P_r, (1, 1) \rangle \not\models \alpha$, so $\langle P_r, (r+1, r) \rangle \not\models C\alpha$. On the other hand, it is easily seen that

$$\langle P_r, (r+1, r) \rangle \models K_1 \underbrace{K_2 K_1 K_2 K_1 \dots K_2 K_1}_{2r-2} \alpha$$

It follows that for all $i_1, i_2, \dots, i_{2r-2} \in \{1, 2\}$,

$$\langle P_r, (r+1, r) \rangle \models K_{i_1} \dots K_{i_{2r-2}} \alpha,$$

showing that the bound in Theorem 3.3 cannot be improved.

4 Alternate Definitions of Knowledge

According to Halpern and Moses, "If Cp is to be attained, all processors must start supporting it simultaneously."³ Unfortunately the notion of two distant events occurring simultaneously has no meaning in modern physics. Do Halpern and Moses plus Einstein imply that no real distributed system ever achieves new common knowledge? A corollary would be that no real, synchronous distributed system can exist.

A look at our example protocols A and B reveals that they are realistic. Recall that new common knowledge among the n processors is attainable in B but not in A . This is all the more confusing because in a very strong sense A and B are isomorphic protocols (cf. [CM85]).⁴

The difference between protocols A and B concerns the granularity at which processors in the two protocols may introspect. In B , processors are only allowed

³[HM84], Lemma 2.

⁴We will call a pair of protocols such as A and B , all of whose interactions are accomplished by a series of messages, *isomorphic* if the set of messages sequences they generate is identical up to permutations which do not switch the order of a send and a receive by the same participant, nor the order of a send and its corresponding receive.

to think about what they know at the start of each write phase. When two isomorphic structures differ on some property, we become very suspicious about whether or not that property is well defined. In the present case we must reexamine our definitions of knowledge and common knowledge.

Let \mathcal{P} be any protocol and let $S \subseteq R_{\mathcal{P}}$ be any subset of reachable global states. For each i , let \sim_S^i be the restriction of \sim^i to $S \times S$. We can now generalize our previous definition of neighborhood. Let $G \subseteq \{1, \dots, n\}$ be a group of the participants in a protocol. For any $p \in S$, define the G -neighborhood of p with respect to S as follows:

$$N_G^S(p) = \{q \mid (\exists r \geq 0)(\exists i_1, \dots, i_r \in G)(\exists p_1, \dots, p_{r-1} \in S) \\ [p \sim_{S}^{i_1} p_1 \sim_{S}^{i_2} p_2 \dots p_{r-1} \sim_{S}^{i_r} q]\}$$

Intuitively, i knows only about the global states in S . We thus define $K_i^S \alpha$ as follows: for $p \in S$,

$$\langle \mathcal{P}, p \rangle \models K_i^S \alpha \equiv \forall q \in N_i^S(p) (\langle \mathcal{P}, q \rangle \models \alpha).$$

Similarly, we can define common knowledge in G with respect to S :

$$\langle \mathcal{P}, p \rangle \models C_G^S \alpha \equiv \forall q \in N_G^S(p) (\langle \mathcal{P}, q \rangle \models \alpha).$$

It is easy to see that the following generalization of Theorem 2.2 holds:

Theorem 4.1 *Let $S \subseteq R_{\mathcal{P}}$ and let $G \subseteq \{1, \dots, n\}$. For $p \in S$ the following two statements are equivalent:*

1. $\langle \mathcal{P}, p \rangle \models C_G^S \alpha$.
2. $(\forall r \geq 0)(\forall i_1, \dots, i_r \in G) (\langle \mathcal{P}, p \rangle \models K_{i_1}^S K_{i_2}^S \dots K_{i_r}^S \alpha)$.

The following theorem shows that for any protocol \mathcal{P} and any nonempty $S \subseteq R_{\mathcal{P}}$, the operators K^S and C^S satisfy the standard S5 axioms for knowledge operators. It follows that if we consider the protocol \mathcal{A} with $S = R_{\mathcal{A}}$, then we get a quite reasonable definition of knowledge and common knowledge for which the asynchronous protocol \mathcal{A} does attain new common knowledge. This contradicts Theorem 3 of [HM84]. More importantly, these observations show the definitions of knowledge and common knowledge needed to make useful progress in the understanding of distributed protocols are much more subtle than one might have at first thought.

Theorem 4.2 *For any protocol \mathcal{P} , any nonempty $S \subseteq R_{\mathcal{P}}$, and $G \subseteq \{1, \dots, n\}$, the operators K_i^S and C_G^S satisfy the standard S5 axioms for modal operators.*

Proof This is immediate from the fact that each \sim_S^i is an equivalence relation. ■

5 Conclusions

We have given precise formulations of distributed protocols. For any subset S of the reachable states, we have given a precise definition of knowledge and common knowledge with respect to S . We have presented theorems outlining some cases where new common knowledge can be attained and some cases where it cannot. Most strikingly, we have shown that in some situations two plausible choices for S can give completely different results.

One can now ask the question, "For which sets of protocols is there a 'best' choice for S ?" and thus a 'best' definition for knowledge and common knowledge. We suspect that in at least certain situations there may be such a best S , and that in this case knowledge and common knowledge with respect to S may be valuable tools.

Many arguments in distributed systems are first formulated at the intuitive level of what certain processors 'know' at certain points in the computation. With precise definitions for these concepts, it may be easier to formulate clear and correct proofs. We believe that considerable work is needed in order to develop logical tools and demonstrate their usefulness on problems of interest in distributed systems.

6 References

References

- [CM85] K. Mani Chandy and Jayadev Misra, "How Processes Learn," *Fourth ACM Symp. on Principles of Distributed Computing* (1985), 204-214.
- [HM84] J. Y. Halpern and Y. Moses, "Knowledge and Common Knowledge in a Distributed Environment," *Third ACM Symp. on Principles of Distributed Computing* (1984), 50-61.
- [LF81] Nancy A. Lynch and Michael J. Fischer, "On Describing the Behavior and Implementation of Distributed Systems," *Theoretical Comp. Sci.* 19 (1981), 17-43.

DISTRIBUTION LIST

Office of Naval Research Contract N00014-82-K-0154

Michael J. Fischer, Principal Investigator

Defense Technical Information Center
Building 5, Cameron Station
Alexandria, VA 22314
(12 copies)

Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Dr. R.B. Grafton, Scientific
Officer (1 copy)

Information Systems Program (437)
(2 copies)

Code 200 (1 copy)
Code 455 (1 copy)
Code 458 (1 copy)

Office of Naval Research
Branch Office, Pasadena
1030 East Green Street
Pasadena, CA 91106
(1 copy)

Naval Research Laboratory
Technical Information Division
Code 2627
Washington, D.C. 20375
(1 copy)

Dr. A.L. Slafkosky
Scientific Advisor
Commandant of the Marine Corps
Code RD-1
Washington, D.C. 20380
(1 copy)

Naval Ocean Systems Center
Advanced Software Technology Division
Code 5200
San Diego, CA 92152
(1 copy)

Mr. E.R. Gleissner
Naval Ship Research and Development Center
Computation and Mathematics Department
Bethesda, MD 20084
(1 copy)

Captain Grace M. Hopper
Naval Data Automation Command
Washington Navy Yard
Building 166
Washington, D.C. 20374
(1 copy)

Defense Advance Research Projects Agency
ATTN: Program Management/MIS
1400 Wilson Boulevard
Arlington, VA 22209
(3 copies)

END

FILMED

12-85

DTIC